Journal of Banking & Finance 37 (2013) 4381–4403

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

# Forecasting the return distribution using high-frequency volatility measures

Jian Hua, Sebastiano Manzan *

*Department of Economics & Finance, Zicklin School of Business, Baruch College, CUNY, United States*

ARTICLE INFO

*Article history:*
Received 30 November 2012
Accepted 4 August 2013
Available online 17 August 2013

*JEL Classification:*
C14
C22
C53

*Keywords:*
Realized volatility
Quantile regression
Density forecast
Value-at-risk

ABSTRACT

The aim of this paper is to forecast (out-of-sample) the distribution of financial returns based on realized volatility measures constructed from high-frequency returns. We adopt a semi-parametric model for the distribution by assuming that the return quantiles depend on the realized measures and evaluate the distribution, quantile and interval forecasts of the quantile model in comparison to a benchmark GARCH model. The results suggest that the model outperforms an asymmetric GARCH specification when applied to the S&P 500 futures returns, in particular on the right tail of the distribution. However, the model provides similar accuracy to a GARCH (1,1) model when the 30-year Treasury bond futures return is considered.



## 1. Introduction

Until recently, the predominant approach in modeling the conditional distribution of returns was represented by the ARCH-GARCH model proposed by Engle (1982) and Bollerslev (1986) and followed by a myriad of sophisticated refinements to the baseline model. The GARCH model introduces time variation in the conditional distribution largely through the conditional variance, and has been successful in explaining several empirical features of asset returns, such as fat tails and the slowly decaying autocorrelation in squared returns. While the GARCH model assumes a parametric form for the latent variance of returns, the recent availability of high-frequency data has sparked a growing literature of volatility estimators that do not require researchers to specify a model. The so-called realized volatility literature (see Andersen and Bollerslev (1998), Andersen et al. (2001a,b), and among others) uses high-frequency data to proxy for the volatility of lower frequency returns, for instance, summing intraday squared returns to estimate the daily variance. In this way, the latent variance process is observable and measured by realized volatility which facilitates the task of modeling and forecasting using time series models. Several recent papers incorporate these measures within a parametric volatility model for the dynamics of daily returns (see Shephard and Sheppard, 2010; Brownlees and Gallo, 2010; Maheu and McCurdy, 2011; Hansen et al., 2012).

In this paper we propose to relate the realized volatility measures and returns by assuming that these measures represent the driving force for the variation of the quantiles of the cumulative multi-period return distribution. In particular, the flexibility of the quantile regression model (see Koenker and Bassett, 1978) allows to consider several specifications that include smoothed versions of the realized volatility measures, the return standardized by realized volatility and nonlinear transformations of the return that are considered to account for the leverage effect. The fact that the parameters of the quantile regression model are specific to each quantile level allows the variables to have heterogeneous effects in different parts of the return distribution. In addition, the quantile model does not require to specify a distribution for the error as it is instead the case for GARCH models or for models based on realized volatility measures. Hence, the flexibility in choosing the most appropriate explanatory variables, the adaptability of the effect of these variables at each quantile level, and the distribution-free character of the method are the three characteristics that distinguish our approach from the models recently proposed in the literature that relate realized volatility and returns. The application of quantile regression to modeling and forecasting financial returns has experienced a recent surge of interest due to

* Corresponding author. Address: Department of Economics and Finance, Baruch College, 55 Lexington Avenue, New York, NY 10010, USA. Tel.: +1 646 312 3408.

*E-mail address:* sebastiano.manzan@baruch.cuny.edu (S. Manzan).

0378-4266/$ - see front matter 
http://dx.doi.org/10.1016/j.jbankfin.2013.08.002

the emergence of risk management and its focus on forecasting the return quantiles (see Engle and Manganelli, 2004; Xiao and Koenker, 2009; Zikes, 2010; Gaglianone et al., 2011). Another aspect that distinguishes our paper is the method adopted to evaluate the performance of quantile and distribution forecasts. A common loss function used in the comparison of density forecasts is the logarithmic score rule, which rewards forecasts that have higher density at the realization of the variable being forecast (see, among others, Bao et al. (2007), Amisano and Giacomini (2007), Maheu and McCurdy (2011), and Shephard and Sheppard (2010) for two applications in the realized volatility literature). Although this is certainly a relevant criterion to consider, it does not reward forecasts that assign high probabilities to values close to the realization in addition to the fact that it cannot be easily adapted to evaluate specific areas of the distribution, for instance the left or right tail. Gneiting and Raftery (2007) and Gneiting and Ranjan (2011) discuss alternative rules that overcome these problems, and we consider several of these rules to evaluate different characteristics of the return distribution. In particular, we use the Quantile Score rule represented by the tick loss function which is targeted to quantile forecasts, such as VaR (e.g. Clements et al., 2008). Instead of focusing on few quantiles of interests, we examine several of them that span the complete return distribution which allows to evaluate the forecast performance of the competing models in different areas of the distribution. Furthermore, we also consider a weighted version of the Quantile Score rule that evaluates specific areas of the forecast distribution, for instance the left and right tail or the center of the distribution, and a scoring rule that evaluates interval forecasts at the 50% and 90% level.

In the empirical application we consider the S&P 500 index futures (SP) and the 30-year Treasury bond futures (US) and forecast out-of-sample the cumulative return at the 1, 2, and 5-day horizons. We evaluate and compare the forecasts from the realized volatility quantile model to those of a benchmark GARCH model represented by the GJR specification (Glosten et al., 1993) for the SP returns and the simple GARCH (1,1) for the US returns. We consider several high-frequency measures of volatility that have been proposed in the literature, including several adjustments that account for the presence of microstructure noise and jumps. The results for the SP returns indicate that the distribution forecasts at the 1-day horizon from the realized volatility models outperform those from the GJR model, with the improved performance mostly deriving from the better ability to forecast the right tail of the return distribution. Only the specifications that include an asymmetric effect are able to beat GJR in modeling the left tail, and significantly so for quantile levels between 20% and 30%. Furthermore, the comparison suggests that the realized measures of volatility considered deliver very similar results, thus indicating that filtering out the effect of jumps and microstructure noise does not improve the (out-of-sample) forecasting ability in any part of the return distribution. In addition, we also consider some quantile specifications that use absolute daily returns (and their transformations), instead of the realized measures, and the evidence indicates that their forecasts do not outperform those from the GJR benchmark. This result indicates that the flexibility of the quantile model combined with the (absolute) returns produces forecasts that have similar accuracy relative to GARCH models, although it does not require to assume a parametric specification. In addition, the realized volatility measures provide valuable information that can be used to improve the accuracy of forecasts relative to (quantile or GARCH) models that only use returns. However, the evidence for the US bond return shows that the realized volatility quantile models provide similarly accurate forecasts relative to those of the benchmark GARCH (1,1) model at all horizons. In this case thus the realized volatility measures do not provide additional forecasting power for the return distribution compared to what is already embedded in daily returns, contrary to the results for the equity index returns.

The paper is organized in this manner. Section 2 describes the realized measures of volatility that are considered in this paper, while Section 3 introduces the GARCH specifications and the semi-parametric model that we propose to incorporate the realized measures in modeling return quantiles. Section 4 describes the forecast evaluation methods and Section 5 reports the results of the empirical application. Finally, Section 7 concludes.

## 2. Realized volatility estimators

The availability of high-frequency data has sparked the development of methods to estimate the (latent) volatility of financial returns that do not require the specification of a model. The most well-known quantity is realized volatility which is obtained by summing intra-day squared returns and can be used to proxy for integrated volatility (see Andersen and Bollerslev, 1998; Barndorff-Nielsen and Shephard, 2002a,b; Meddahi, 2002). In this section, we present several realized volatility measures that are later used in our empirical application.

Denote the intra-day return in day $t$ by $r_{t,i} = ln(P_{t_i}) - ln(P_{t_{i-1}})$, where $i = 1, 2, \ldots, m$ indicates the intra-day interval and $P_{t_i}$ the asset price in interval $i$ of day $t$. The realized volatility estimator in day $t$, denoted by $RV_t$, represents a model-free estimator of the daily quadratic variation at sampling frequency $m$ and is given by

$$RV_t = \sum_{i=1}^{m} r_{t,i}^2. \tag{1}$$

The asymptotic distribution of $RV_t$ has been studied by Andersen and Bollerslev (1998), Andersen et al. (2001b, 2003), Barndorff-Nielsen and Shephard (2002a), and among others. An important role in the construction of the measure is played by the selection of the sampling frequency $m$ which is complicated by several market microstructure issues (see, e.g. Ait-Sahalia et al. (2005a,b), Bandi and Russell (2008), Hansen and Lunde (2006b), Barndorff-Nielsen et al. (2008), and among others). In our empirical application, we use a five-minute sampling frequency, which has been shown in the literature to strike a reasonable balance between the desire for as finely sampled observations as possible and robustness to market microstructure contaminations.

Despite the careful selection of sampling frequency, market-microstructure dynamics could still cause $RV_t$ to be a biased and inconsistent estimator of volatility. Thus, we also consider estimators with adjustments that reduce market microstructure frictions present in high-frequency returns. We adopt a kernel-based estimator of realized volatility suggested by Hansen and Lunde (2006b), which employs Bartlett weights,

$$RV(q)_t = \sum_{i=1}^{m} r_{t,i}^2 + 2\sum_{w=1}^{q}\left(1 - \frac{w}{q+1}\right)\sum_{i=1}^{m-w} r_{t,i} r_{t,i+w}, \tag{2}$$

where $r_{t,i}$ is defined as above and $\left(1 - \frac{w}{q+1}\right)$ represents the weight that follows a Bartlett scheme. This estimator utilizes higher-order auto-covariances to eliminate the bias of $RV_t$, and is also guaranteed to be non-negative. The asymptotic properties of the estimator are discussed by Barndorff-Nielsen et al. (2008).

Volatility can also experience frequent jumps. Andersen et al. (2007) suggest that most of the predictable variation in the volatility stems from the continuous price path variability while the predictability of the jump component of volatility is typically minor. However, Wright and Zhou (2009) find that measures of realized jumps are useful predictors for bond risk premia. We follow the